arXiv:1603.02443v2 [stat.ML] 9 Mar 2016

Note on the equivalence of hierarchical variational models and auxiliary deep generative models

Niko Brümmer
AGNITIO Research South Africa

Strand, March 2016

Abstract

This note compares two recently published machine learning methods for constructing flexible, but tractable families of variational hidden-variable posteriors. The first method, called *hierarchical variational models* enriches the inference model with an extra variable, while the other, called *auxiliary deep generative models*, enriches the generative model instead. We conclude that the two methods are mathematically equivalent.

1 Introduction

In machine learning, there is an ongoing revival of the use of variational Bayes (VB) to deal with complex probabilistic models with hidden variables. The revival is driven by the use of stochastic methods to approximate the VB lower bound and associated gradients. See for example [1, 2, 5]. The advantages include automated inference [3] and also that they are applicable to a much wider class of probabilistic models. While the basic recipes are limited in the flexibility of the approximate hidden-variable posteriors, there are ongoing efforts to make them more flexible. For example, [4] allows complex reparametrizations of the hidden variables. These reparametrizations are however still limited by an invertibility constraint and by requiring computation of Jacobian determinants. This note summarizes and compares further progress in this regard as published in [6, 7, 8].

2 Problem statement

We start by defining of the kind of probabilistic models of interest, introduce the stochastic VB solution to deal with them and state a limitation of stochastic VB w.r.t. the approximate posteriors.

Our model has the following form. Let $\mathbf{x}$ be observed and $\mathbf{z}$ be hidden. The generative model is

$$P(\mathbf{x}, \mathbf{z}) = P(\mathbf{x} \mid \mathbf{z})P(\mathbf{z}) \tag{1}$$

where we assume we can evaluate both factors in the RHS and therefore also the LHS. However, the marginal $P(\mathbf{x})$ and posterior $P(\mathbf{z}|\mathbf{x})$ are assumed intractable. The model, $P(\mathbf{x}, \mathbf{z})$, will usually be conditioned on some parameters, but these parameters play no role in this discussion and are considered implicit in the definition of P.

Variational Bayes (VB) allows us to employ some convenient family of parametrized distributions, $Q(\mathbf{z}|\theta)$, where θ can be chosen to approximate the intractable posterior:

$$Q(\mathbf{z} \mid \theta) \approx P(\mathbf{z} \mid \mathbf{x}) \tag{2}$$

In the standard VB recipe, θ is found by maximizing the evidence lower bound (ELBO):

$$\mathcal{L}(\theta) = \int Q(\mathbf{z} \mid \theta) \log \frac{P(\mathbf{x}, \mathbf{z})}{Q(\mathbf{z} \mid \theta)} \, d\mathbf{z} \tag{3}$$

If the ELBO integral cannot be evaluated in closed form, it can still be approximated stochastically by using, say N, samples from $Q(\mathbf{z}|\theta)$:

$$\mathcal{L}(\theta) \approx \frac{1}{N} \sum_{\mathbf{z} \sim Q(\mathbf{z}|\theta)} \log \frac{P(\mathbf{x}, \mathbf{z})}{Q(\mathbf{z} \mid \theta)} \tag{4}$$

The gradients w.r.t. θ, which are needed for the optimization, require some further tricks, for example the reparametrization trick [1]. See [5] for a review of methods to deal with the gradients.

The above VB recipe can only be applied if we can evaluate $Q(\mathbf{z}|\theta)$. Below we summarize two published methods, *hierarchical variational models* [6, 7] and *auxiliary deep generative models* [8] that construct approximate posteriors that cannot be evaluated, but which can nevertheless be used in slightly more complex VB strategies.

3 Hierarchical Variational Models

This is a summary and discussion of ideas used in [6, 7]. The VB posterior is assembled by involving another hidden variable, say $\boldsymbol{\lambda}$, which must be marginalized out because it plays no role in the generative model $P(\mathbf{x}, \mathbf{z})$. The full inference model is defined as:

$$Q(\mathbf{z}, \boldsymbol{\lambda} \mid \theta) = Q(\boldsymbol{\lambda} \mid \theta) Q(\mathbf{z} \mid \boldsymbol{\lambda}, \theta) \tag{5}$$

We assume that we can evaluate and sample from both factors in the RHS. This means the joint density (LHS) can also be evaluated and sampled from—i.e. by doing ancestral sampling. The marginal

$$Q(\mathbf{z} \mid \theta) = \int Q(\mathbf{z}, \boldsymbol{\lambda} \mid \theta)\, d\boldsymbol{\lambda} \tag{6}$$

is however assumed to be intractable—and by implication also the posterior $Q(\boldsymbol{\lambda}|\mathbf{z}, \theta)$.

We deal with this inference model via a nested VB recipe, where we introduce yet another approximate posterior:

$$R(\boldsymbol{\lambda} \mid \mathbf{z}, \phi) \approx Q(\boldsymbol{\lambda} \mid \mathbf{z}, \theta) \tag{7}$$

that is chosen from a family of distributions that we can evaluate.[1] We now modify the ELBO by subtracting a non-negative term, namely an expected KL divergence, so that the *modified ELBO* is still a lower bound to the evidence:

$$\mathcal{L}(\theta, \phi) = \mathcal{L}(\theta) - \int Q(\mathbf{z} \mid \theta) \left[\int Q(\boldsymbol{\lambda} \mid \mathbf{z}, \theta) \log \frac{Q(\boldsymbol{\lambda} \mid \mathbf{z}, \theta)}{R(\boldsymbol{\lambda} \mid \mathbf{z}, \phi)}\, d\boldsymbol{\lambda} \right] d\mathbf{z} \tag{8}$$

where the factor in square brackets is the KL-divergence. If we maximize $\mathcal{L}(\theta, \phi)$ w.r.t. ϕ, it will be forced closer to the original $\mathcal{L}(\theta)$, reaching equality if $R(\boldsymbol{\lambda}|\mathbf{z}, \phi) = Q(\boldsymbol{\lambda}|\mathbf{z}, \theta)$ for every value of z in the support of $Q(\mathbf{z}|\theta)$. By this construction we now have:

$$\log P(\mathbf{x}) \geq \mathcal{L}(\theta) \geq \mathcal{L}(\theta, \phi). \tag{9}$$

We now re-arrange $\mathcal{L}(\theta, \phi)$ in terms of tractable components. Since we can sample from $Q(\mathbf{z}, \boldsymbol{\lambda}|\theta)$, we are happy with expectations formed over $\mathbf{z}, \boldsymbol{\lambda}$. So we can rewrite:

$$\mathcal{L}(\theta, \phi) = \mathcal{L}(\theta) + \int \int Q(\mathbf{z}, \boldsymbol{\lambda} \mid \theta) \log \frac{R(\boldsymbol{\lambda} \mid \mathbf{z}, \phi)}{Q(\boldsymbol{\lambda} \mid \mathbf{z}, \theta)}\, d\boldsymbol{\lambda}\, d\mathbf{z} \tag{10}$$

[1] We don't need to be able to sample from it.

Since the integrand in (3) is not a function of λ, it can be rewritten as:

$$\mathcal{L}(\theta) = \int \int Q(\mathbf{z}, \boldsymbol{\lambda} \mid \theta) \log \frac{P(\mathbf{x}, \mathbf{z})}{Q(\mathbf{z} \mid \theta)} \, d\boldsymbol{\lambda} \, d\mathbf{z} \tag{11}$$

Both (10) and (11) have intractable factors, namely $Q(\boldsymbol{\lambda}|\mathbf{z}, \theta)$ and $Q(\mathbf{z}|\theta)$. But when we combine them, these factors are multiplied to form the joint distribution, $Q(\mathbf{z}, \boldsymbol{\lambda} \mid \theta)$, which we *can* evaluate. This gives:

$$\begin{aligned}
\mathcal{L}(\theta, \phi) &= \int \int Q(\mathbf{z}, \boldsymbol{\lambda} \mid \theta) \left[\log \frac{P(\mathbf{x}, \mathbf{z})}{Q(\mathbf{z} \mid \theta)} + \log \frac{R(\boldsymbol{\lambda} \mid \mathbf{z}, \phi)}{Q(\boldsymbol{\lambda} \mid \mathbf{z}, \theta)} \right] d\boldsymbol{\lambda} \, d\mathbf{z} \\
&= \int \int Q(\mathbf{z}, \boldsymbol{\lambda} \mid \theta) \log \frac{P(\mathbf{x}, \mathbf{z}) R(\boldsymbol{\lambda} \mid \mathbf{z}, \phi)}{Q(\mathbf{z}, \boldsymbol{\lambda} \mid \theta)} \, d\boldsymbol{\lambda} \, d\mathbf{z} \\
&\approx \frac{1}{N} \sum_{\mathbf{z}, \boldsymbol{\lambda} \sim Q(\mathbf{z}, \boldsymbol{\lambda} | \theta)} \log \frac{P(\mathbf{x}, \mathbf{z}) R(\boldsymbol{\lambda} \mid \mathbf{z}, \phi)}{Q(\mathbf{z}, \boldsymbol{\lambda} \mid \theta)}
\end{aligned} \tag{12}$$

3.1 Example

We can use this recipe to involve a general-purpose feed-forward neural net to define a flexible family for the variational posterior. Let $\mathbf{z} = (z_1, \ldots, z_n) \in \mathbb{R}^n$ and let $\boldsymbol{\lambda} \in \mathbb{R}^m$. Let:

$$Q(\mathbf{z} \mid \boldsymbol{\lambda}, \theta) = \prod_{i=1}^{n} \mathcal{N}\big(z_i \mid \mu_i(\boldsymbol{\lambda}; \theta), \sigma_i^2(\boldsymbol{\lambda}; \theta)\big) \tag{13}$$

where we employ a neural net, parametrized by θ, to map $\boldsymbol{\lambda}$ to the means and variances of a diagonal Gaussian distribution for $\mathbf{z}$. Let

$$Q(\boldsymbol{\lambda} \mid \theta) = Q(\boldsymbol{\lambda}) = \mathcal{N}(\boldsymbol{\lambda} \mid \mathbf{0}, \mathbf{I}) \tag{14}$$

This gives a very flexible continuous *Gaussian mixture*:

$$Q(\mathbf{z} \mid \theta) = \int \mathcal{N}(\boldsymbol{\lambda} \mid \mathbf{0}, \mathbf{I}) \prod_{i=1}^{n} \mathcal{N}\big(z_i \mid \mu_i(\boldsymbol{\lambda}; \theta), \sigma_i^2(\boldsymbol{\lambda}; \theta)\big) \, d\boldsymbol{\lambda} \tag{15}$$

We can sample from it using ancestral sampling, but we cannot evaluate the marginal density $Q(\mathbf{z} \mid \theta)$. This makes the above recipe applicable.

We can form $R(\boldsymbol{\lambda}|\mathbf{z}, \phi)$ similarly with a neural net, parametrized by ϕ, that maps $\mathbf{z}$ to the parameters of a Gaussian distribution for $\boldsymbol{\lambda}$.

4 Auxiliary Deep Generative Models

We now summarize the other method, published in [8] and show that it is very close to the method in [6, 7]. Let's recall (12):

$$\mathcal{L}(\theta, \phi) = \int \int Q(\mathbf{z}, \boldsymbol{\lambda} \mid \theta) \log \frac{P(\mathbf{x}, \mathbf{z}) R(\boldsymbol{\lambda} \mid \mathbf{z}, \phi)}{Q(\mathbf{z}, \boldsymbol{\lambda} \mid \theta)} \, d\boldsymbol{\lambda} \, d\mathbf{z} \tag{16}$$

where P belongs to the generative model and Q, R to the inference model. However, nothing changes mathematically if we rename R to P and thereby change the role of λ to be a hidden variable in both generative and inference models:

$$\begin{aligned} \mathcal{L}(\theta, \phi) &= \int \int Q(\mathbf{z}, \boldsymbol{\lambda} \mid \theta) \log \frac{P(\mathbf{x}, \mathbf{z}) P(\boldsymbol{\lambda} \mid \mathbf{z}, \phi)}{Q(\mathbf{z}, \boldsymbol{\lambda} \mid \theta)} \, d\boldsymbol{\lambda} \, d\mathbf{z} \\ &= \int \int Q(\mathbf{z}, \boldsymbol{\lambda} \mid \theta) \log \frac{P(\mathbf{x}, \mathbf{z}, \boldsymbol{\lambda} \mid \phi)}{Q(\mathbf{z}, \boldsymbol{\lambda} \mid \theta)} \, d\boldsymbol{\lambda} \, d\mathbf{z} \end{aligned} \tag{17}$$

which is just the standard ELBO for an *extended model*, $P(\mathbf{x}, \mathbf{z}, \boldsymbol{\lambda}|\phi)$, with two hidden variables, $\mathbf{z}$ and $\boldsymbol{\lambda}$. If the original model, $P(\mathbf{x}, \mathbf{z}) = P(\mathbf{x}|\mathbf{z})P(\mathbf{z})$ is visualized as:

Then the extended model implied by (17), is

$$P(\mathbf{x}, \mathbf{z}, \boldsymbol{\lambda}|\phi) = P(\mathbf{x} \mid \mathbf{z}) P(\mathbf{z}) P(\boldsymbol{\lambda} \mid \mathbf{z}, \phi)$$

and can be visualized as:

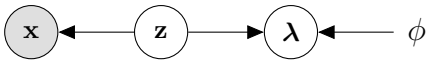

Since $\boldsymbol{\lambda}$ is hidden, the original variables $\mathbf{z}$ and $\mathbf{x}$ are independent of ϕ, so that ϕ may be modified at will, without changing the generative description of the data: $P(\mathbf{x}|\phi) = P(\mathbf{x})$. As the analysis in the previous section showed, by maximizing (17) w.r.t. ϕ, we are instead effectively improving the *inference* model.

As pointed out in both [7] and [8], it turns out this recipe still works for a slightly more complex extended model, visualized as:

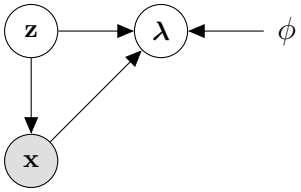

That is, we have replaced $R(\boldsymbol{\lambda}|\mathbf{z}, \phi)$ by $R(\boldsymbol{\lambda}|\mathbf{z}, \mathbf{x}, \phi)$. It should be clear from (8) that this can be done without changing the relationship (9) between the standard and modified lower bounds.

5 Conclusion

Mathematically, the two methods are the same. Perhaps the derivation of the method is easier to understand from the *hierarchical variational model* viewpoint, while software implementation could be facilitated by the *auxiliary deep generative models* viewpoint.

References

[1] Diederik P Kingma and Max Welling"Auto-Encoding Variational Bayes", December 2013. http://arxiv.org/abs/1312.6114.

[2] Danilo Jimenez Rezende, Shakir Mohamed and Daan Wierstra, "Stochastic backpropagation and approximate inference in deep generative models", January 2014. http://arXiv:1401.4082,2014.

[3] Alp Kucukelbir, Dustin Tran, Rajesh Ranganath, Andrew Gelman and David M. Blei, "Automatic Differentiation Variational Inference", March 2016. http://arxiv.org/abs/1603.00788.

[4] Danilo Jimenez Rezende, Shakir Mohamed, "Variational Inference with Normalizing Flows", May 2015. http://arxiv.org/abs/1505.05770.

[5] Michalis K. Titsias, "Local Expectation Gradients for Doubly Stochastic Variational Inference", March 2015/ http://arxiv.org/abs/1503.01494.

[6] Rajesh Ranganath, Dustin Tran and David M. Blei, "Hierarchical Variational Models", November 2015. http://arxiv.org/abs/1511.02386.

[7] Dustin Tran, Rajesh Ranganath and David M. Blei, "Variational Gaussian Process", November 2015. http://arxiv.org/abs/1511.06499.

[8] Lars Maaløe, Casper Kaae Sønderby, Søren Kaae Sønderby, Ole Winther, "Auxiliary Deep Generative Models", February 2016. http://arxiv.org/abs/1602.05473.